Enhanced Product Generation at NASA Data Centers Through Grid Technology

Bruce R. Barkstrom+, Thomas H. Hinke*, Shradha Gavali* and William J. Seufzer!

* NASA Ames Research Center, + NASA Langley Research Center, !Science Applications International Corporation, Hampton, VA

Abstract

This paper describes how grid technology can support the ability of NASA data centers to provide customized data products. A combination of grid technology and commodity processors are proposed to provide the bandwidth necessary to perform customized processing of data, with customized data subsetting providing the initial example. This customized subsetting engine can be used to support a new type of subsetting, called phenomena-based subsetting, where data is subsetted based on its association with some phenomena, such as mesoscale convective systems or hurricanes. This concept is expanded to allow the phenomena to be detected in one type of data, with the subsetting requirements transmitted to the subsetting engine to subset a different type of data. The subsetting requirements are generated by a data mining system and transmitted to the subsetter in the form of an XML feature index that describes the spatial and temporal extent of the phenomena. For this work, a grid-based mining system called the Grid Miner is used to identify the phenomena and generate the feature index. This paper discusses the value of grid technology in facilitating the development of a high performance customized product processing and the coupling of a grid mining system to support phenomena-based subsetting.

1. Introduction

NASA has a number of archives that hold large amounts of data generated by various satellites. One group of archives is dedicated to the storage of Earth science data, with the nine archives associated with the Earth Observing System Data and Information System (EOSDIS) holding the bulk of the NASA's Earth science data. Currently users can access one of the EOSDIS archives through various data gateways spread around the world that are all listed on

http://redhook.gsfc.nasa.gov/~imswww/pub/imswelcome/imswwwsites.html. Through these gateways, the user can perform a search for data across all of the EOSDIS archives, based on the name of the instrument that was used to capture the data (e.g., TRMM TMI - Tropical Rainfall Measuring Mission Microwave Imager or CERES – Cloud and Earth's Radiant Energy System), various keywords such as parameters (e.g., Cloud Liquid Water/Ice), processing level, and numerous others. The result of a search is a list of data sets that satisfy the search criteria. From this list the user can select a particular data set and view the file that can be ordered for that data set. The smallest orderable quantity of data is called a granule and may cover a fraction of an orbit, a single orbit, a

single day or some longer period. Under current practice, the user can place his order though the data gateway's web page or through web interfaces maintained by the individual data centers. If only a small amount of data is ordered (say, a few gigabytes or less) then the data will be pulled from the archive in a few hours and placed on an FTP server. If a larger amount of data is ordered, then the data will be shipped to the user on some media (such as digital tape).

This paper is investigating technology that will improve this process by allowing the user to specify some custom processing to be performed at the data center that will transform the data into a form that is more useable to him or her. Under this approach, rather than getting the standard files of data, and then having to perform custom processing at the user's site, the user can request that the data center produce some custom products. In the simplest case, this custom product can involve subsetting the data. For example, each granule might cover an entire orbit or an entire day (e.g., 14 orbits for some satellites). while the user is interested only in data that covers California. Alternatively, the user may be interested in time series data that covers California over the last three years. If the science team producing the data has not already done the selection, this interest could require extensive processing to go though three years of data and gather only the data that covers California. The user may also want only data that is associated with a mesoscale convective system (severe storm), a hurricane or a volcanic eruption. This involves phenomena-based subsetting which requires that a significant amount of processing be expended by the data center if these types of customized product requests are to be handled in the future.

This paper describes research and development work to improve the order production process by performing custom subsetting of the data prior to delivery to the requestor through the user of grid technology to support parallel subsetting of data. The data center portion of this work is being performed at the NASA Langley Research Center Atmospheric Sciences Data Center, which is one of the EOSDIS archives. The data mining portion of this work to support phenomena-based subsetting is being performed by the NASA Ames Research Center.

The next section of this paper will describe the work at the Atmospheric Sciences Data Center to use a combination of grid technologies and commodity processors to provide a high performance engine to perform customized product processing, with subsetting being the initial application. Section 3 will present the approach being used to tie a Grid Miner to the data center's subsetter to support phenomena-based subsetting. Section 4 will conclude with a discussion of how grid technologies are facilitating the development of this system.

2. Customized Product Processing Through Grid Technology

At present, EOSDIS data is stored in robotic tape archives, with data read rates for a single file of about 10 MBps. This read rate limits the amount of data a user can access in a month to about 26 TB – assuming that there are no glitches and that the user makes no mistakes. Unfortunately, users who want to work with long time series, say five years, need access to several times this amount of data and are likely to have to iterate

with their subsetting algorithms three or four times. With this rate of access, it is nearly impossible for users to manipulate and discover interesting items because it may take several years to produce well-validated subsets. In order to maximize use of the data for this type of user, it will be necessary to massively increase the rate at which data can be run through user programs. As a useful goal, we would like to be able to obtain a throughput rate of about 400 MBps or about 1 PB per month. This would allow the user with 26 TB of data to iterate through the entire data set in about one day (or less), allowing much improved ability to remove processing artifacts and errors.

This throughput can be achieved by transferring the data to a large number of disks, with an appropriate number of CPU's to perform the data filtering and subsetting operations. To move toward this goal, the Atmospheric Sciences Data Center at the NASA Langley Research Center will start by transferring a modest number of files out of the archive for storage on disk. Then, they will provide a simple subsetting program that allows users to interactively build a script of instructions similar to what one might do with a relational database. The intent is to allow this program to create a data structure similar to the rows with fields in a database and then to perform four basic functions on the data:

- Use simple queries on the fields to select rows into the subset
- Calculate simple statistics on the fields in the selected rows
- Visualize the relationships between the fields at high display rates that would allow millions of data points to be plotted on a user's browser in under thirty seconds
- Create transformed variables that are placed in new columns of the in-memory data structure

This approach is expected to deliver useful subsets for various NASA science teams. In addition, we intend to explore the use of commodity computing hardware and highly reliable software using grid computing technology to reduce the overall cost of ownership.

Figure 1 shows the schematic architecture of the Atmospheric Sciences Data Center's grid-based data access and production system. This system will initially support parallel subsetting, but could just as easily support parallel reformatting of the data to transform it from its archived format into a format desired by the data requestor, or perform some other processing to customize the data for the user. It is expected that both internal and external data users will interact with this system through a web-services style of interface. At the same time, it is important for the system to be sufficiently automated that intelligent agents could provide files for ingesting data and for obtaining data through the distribution interface – expected to be push and pull FTP. Internally, both the CPU's and the Storage Nodes should be designed as peer-to-peer daemons to increase system reliability and to improve security.

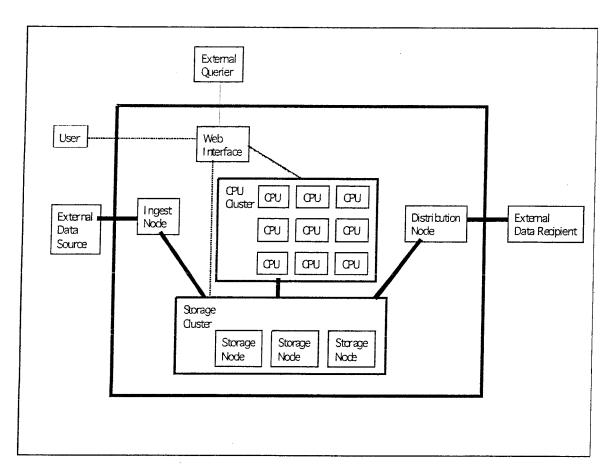


Figure 1. Schematic Diagram for Architecture of Grid-Enabled Data Access and Production System. Heavy lines show high-bandwidth data transfer paths; lighter (dotted) lines show command paths. Both CPU's and Storage Nodes are intended to operate autonomously to increase reliability and security of this architecture.

The NASA Langley Atmospheric Sciences Data Center, has about 500 TB of data currently stored in two systems - both having most of the data on AMASS tape storage systems. Data production uses SGI machines - mainly Origin 2000's or 3800's - with about 150 CPU's in each of the two systems. The original architecture and much of the hardware comes from the early to mid-1990's, when both the data volume and the I/O were perceived as very high in terms of the typical needs of the IT community. With recent developments in hardware and software as well as budgetary pressures, consideration is being given to moving to commodity hardware and open source software including the following:

- Linux for the OS and standard package support
- PostgreSQL for storing and searching our metadata
- Perl and Python for production scripting needs
- gcc for compiler support in Ada95 and Java, Intel ® for C, C++, and FORTRAN
- Hierarchical Data Format (an Open Source format from the HDF Group at NCSA see http://hdf.ncsa.uiuc.edu) and
- NetCDF (another formatting library from Unidata see http://www.unidata.ucar.edu/packages/netcdf/)

- Grid computing tools, notably Storage Resource Broker (SRB) from the San Diego Supercomputing Center (see http://www.npaci.edu/DICE/SRB/)
- Condor-G from the University of Wisconsin (see http://www.cs.wisc.edu/condor/condorg/), and
- Globus toolkit from Argonne National Labs (see http://www.globus.org/toolkit/)

The approach described in this follows the recommendations in the recent report of the National Research Council on Government Data Centers [NRC 2003], which recommends that data centers consider moving from tape storage to disk and incorporating more "bleeding edge" technologies.

While we do not present a detailed architecture here, recent advances in versioning theory for Earth science data products [Barkstrom 2003] will support database use for holding and querying metadata, as well as the ability to perform provenance tracking from "cradle to grave" on the data that will be stored in this system. These features also accord well with the recommendations of the NRC.

The initial application of most interest to the Atmospheric Sciences Data Center is finding ways of subsetting and finding phenomena (like mesoscale convective systems or hurricanes) in the data we have in the Center. From a variety of perspectives, this application suggests the need to be able to stream large files (say 200 MB each) through filtering programs at a rate of about 1 PB/month. This figure would support both migration of large data collections and a moderate number of climatological data users who need reasonable turnaround (say 1 week or less) in working with five or ten years of data records. Assuming that the major bottleneck lies in the I/O, it may mean that we need to do some additional architectural work on how we store the files, but this will await experience with the system being described in this paper.

Perhaps the simplest way of describing the systems currently used by the Atmospheric Sciences Data Center is that they were engineered to deal with data storage - under the assumption that data users wanted only a few files in each access and that the accesses were more or less randomly distributed across all of the files. This access pattern makes for a reasonable match between robotic tape storage, the data storage strategy, and the user access pattern.

In the data access scenario that is driving the new, grid-based subsetting work, users want to stream through data files that span a period of time, perhaps five years or more. Under this scenario, the data center has a need for content based searching with a throughput requirement of approximately 1 PB per month. In some early experiments, the primary bottleneck for this kind of user scenario is in the ability to get data out of the input side of a computational node and through a filtering program in the computer. There is currently no perceived need to worry about parallelizing the individual program code - computation isn't the bottleneck, data bandwidth is. That almost certainly means that the new architecture needs to spread the data out in multiple disks and run it through as many CPU's as we can make available, using coarse-grained parallel processing, with for example, each processor working on data for a different day or orbit.

Since data may be located in different data centers, we anticipate the need to subset related data at multiple data centers. To minimize the amount of data flowing between centers, we would be better off exchanging "pointers" or feature indexes for features discovered in one data center that need to be used to subset data in another data center. This approach should reduce the required network traffic by several orders of magnitude over an approach in which data from multiple centers is all moved to the same location to perform such multi-center, mult-data-set subsetting. It also allows data centers to retain autonomy with respect to their collections at the same time it increases services to users, since the feature indexes that serve as an additional source of metadata to help new user communities find objects of interest can be replicated at both data centers for a very small storage overhead. We expand on this approach in the section that follows.

3. Phenomena-Based Subsetting

Phenomena-based subsetting is a concept that supports the desire to perform research on data from a number of different datasets that are all associated with the same phenomena. In order to support phenomena-based subsetting, the spatial and temporal location of the phenomena of interest must be determined. Since this could involve sifting through a large amount of data to locate phenomena of interest, it represents a potentially good application for data mining, which has been defined as "... the process by which information and knowledge are extracted from a potentially large volume of data using techniques that go beyond a simple search through the data." [Data Mining Workshop 1999] Scientific data mining in general, and Earth Science data mining in particular is characterized by the need to mine possibly large amounts of data that has been captured by satellite-based remote sensors. An example is data from the TMI (TRMM Microwave Imager) instrument, which consists of approximately 230 megabytes (uncompressed) of data per day. Other satellite data can be even more voluminous due to its finer resolution.

This research uses a software system called the Grid Miner [Hinke 2000b] that was developed at the NASA Ames Research Center. The Grid Miner is a grid-enabled version of the stand-alone ADaM data mining system that was developed at the University of Alabama in Huntsville under a NASA research grant [Hinke 2000a, Hinke 1997a]. The Grid Miner is an agent-based mining system in which mining agents are sent to processors on the grid to mine remote data that is accessible from the grid and described in a mining database that has been pre-loaded with the URLs of data to be mined.

Figure 2 shows the architecture that is used to perform phenomena-based subsetting. A user invokes the miner by staging "thin" mining agents to the grid processors that are to support the mining, along with the mining plan (written by the user) that is to guide the mining for the desired phenomena. Based on the mining plan provided, these thin agents are able to grow in capability, through the acquisition of the necessary mining operations required to execute the plan. Each of the mining operations is configured as a shared library executable, with one operation per executable file. As the thin mining agent executes the mining plan, it identifies the operations that are to be used and then uses the grid to transfer the needed shared library executables from a mining operator repository to the grid processors where the mining is to be performed. The use of thin agents

minimizes the size of the agent code that needs to be transferred. This approach of dynamically acquiring needed mining operations means that mining operations could be retrieved from multiple operator repositories, some public, some private and perhaps some for a fee, although this multi-site repository represents future work.

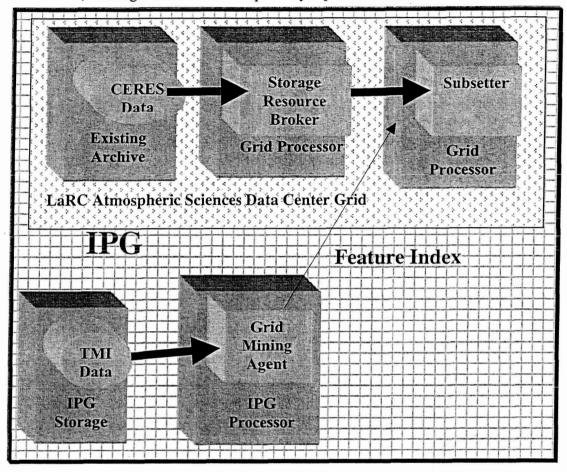


Figure 2. Schematic Diagram for Architecture of Phenomana-based Subsetting. Heavy lines show high-bandwidth data transfer paths; lighter show the path of the XML feature index. Note that this architecture uses both the IPG and an internal Atmospheric Scienes Data Center Grid.

Once the thin agent has grown to have the necessary mining operations to perform the mining plan, the mining agent contacts the mining database to acquire the URLs of the files to be mined. Using the grid, these files are then transferred to the mining site and the mining is performed as specified in the mining plan.

To support phenomena-based subsetting, the Grid Miner will mine the data for the desired phenomena and when found, will circumscribe the phenomena with a convex hull polygon and associated metadata to specify not only the spatial extent of the phenomena, but also its temporal location. These will be output as an XML document in the following form:

<?xml version="1.0" encoding="ISO-8859-1"?>
<polygon_list>

```
<polygon>
 <date time>
 <dateYYYY-MM-DD> 1998-01-01 </dateYYYY-MM-DD>
 <time_of_day in second> 13000</time_of_day in second>
</date time>
<size in_square km> 2050.781250 </size in square km>
<region type> 2</region type>
<vertices>
<number of vertices> 7</number of vertices>
<vertex> <latitude> LATw </latitude> <longitude> LONGx</longitude>
</vertex>
<vertex> <latitude> LATy</latitude> <longitude> LONGz</longitude>
</re>
</polygon>
<polygon>
</polygon>
</polygon list>
```

For the initial phenomena for this work, we are searching for mesoscale convective systems within passive microwave data from the TMI instrument that was obtained from the Goddard Space Flight Center's Distributed Active Archive Center. The mining operation used to search for mesoscale convecive systems was developed at the University of Alabama in Huntsville and uses an algorithm originally suggested in [Devlin 1995]. For the purposes of this experimental work, the TMI data is being stored on storage that is accessible from the NASA Information Power Grid (IPG). The Grid Miner is staged to IPG computational resources, which could be located anywhere on the IPG, with the data to be mined pulled from Ames' storage.

When the mining is completed, the XML document describing the spatial and temporal location of the phenomena of interest will be sent to the subsetting engine. Again, the use of XML for transferring information accords well with the NRC recommendations [NRC 2003]. The XML document includes indexes to the original pixels in the data files that contribute to the object identifications. Such object indices provide the ability to efficiently retrieve the original data belonging to the phenomenon, as well as building metadata for each instance. This approach allows us to develop a database of phenomena instances [Hinke 1997b] that should markedly increase the scientific community's ability to extend the value of its data resources, as suggested several years ago in [Barkstrom 1998].

For this initial work, based on the date and time information provided in the XML document, the appropriate CERES data will be accessed and fed into the subsetter, along with the polygon that describes the spatial and temporal areas that correspond to the mesoscale convective systems from the TMI data. The subsetter will then extract the CERES data that corresponds to the TMI-discovered mesoscale convective system.

It should be noted that in this case (and by intent) the two instruments (TMI and CERES) were both located on the same satellite. Thus, the subset date extracted from the CERES data will have both temporal and spatial congruence with the phenomena discovered in the TMI data. It is not always the case that desired instruments are located on the same satellite, which means that the phenomena detected in data from one satellite may have moved to a slightly different location in the data that is subset. How one would address this problem is beyond the scope of this paper.

4. The Value of Grid Support

Grid technology provides valuable support for the user-centric approach we have described:

- The Storage Resource Broker (SRB) provides a useful way of separating the details of the local storage from the logical structure of the files and directories, reducing the operator overhead associated with storage management, thereby reducing the total cost of ownership.
- The external interfaces to grid architectures, such as that shown in figure 1, provide reliable FTP with the grid-provided single-sign-on environment for connecting the subsetting with access to data stored on a grid-enabled storage system, such as the SRB, increasing the possibilities for fully automated, secure system access, again reducing the total cost of ownership.
- Provides a single-sign-on environment for running the grid miner of available grid resources and handling the transfer of both mining operators and data.
- Provides a single-sign-on environment for injecting feature indexes into the data center's subsetting engine.

5. Acknowledgements

We would like to acknowledge the fact that the Grid Miner is a modification of the ADaM data mining system that was developed by the University of Alabama in Huntsville, under a NASA research grant. While Thomas Hinke, one of the authors of this paper designed a data mining system that was a precursor to the ADaM data mining system and was one of the investigators involved in the development of ADaM, Dr. John Rushing was the principle architect and implementer of the ADaM data mining system. Drs. Sara Graves and Hegge Ranganath also were significant contributors to the development of ADaM. We would also like to acknowledge the valuable comments that William Thigpen and Tony Lisotta at the NASA Ames Research Center made on this paper.

6. References

[Barkstrom 1998] Barkstrom, B. R., "Digital Archive Issues from the Perspective of an Earth Science Data Producer", paper presented at the *Digital Archive Directions (DADS) Workshop*, June 22-26, 1998, available at http://ssdoo.gsfc.nasa.gov/nost/isoas/dads/dads21b.html.

[Barkstrom 2003] Barkstrom, B. R., "Data Product Configuration Management and Versioning in Large-Scale Production of Satellite Scientific Data," in B. Westfechtel and A. van der Hoek (eds.), SCM 2001/2003, Lecture Notes in Computer Science 2649, pp. 118-133, 2003.

[Devlin 1995] Devlin, K.I. Application of the 85 Ghz Ice Scattering Signature to a Global Study of Mesoscale Convective Systems. Master's thesis, Meteorology, Texas A&M University, August 1995.

[Hinke 2000a] Hinke, Thomas H., John Rushing, Heggere Ranganath and Sara J. Graves, "Techniques and Experience in Mining Remotely Sensed Satellite Data," *Artificial Intelligence Review: Issues on the application of data mining*, Vol. 14, No. 6, December 2000, pp. 503-531.

[Hinke 2000b] Hinke, Thomas H. and Jason Novotny, "Data Mining on NASA's Information Power Grid," *Proceedings Ninth IEEE International Symposium on High Performance Distributed Computing*, Pittsburgh, Pennsylvania, August, 2000.

[Hinke 1997a] Hinke, Thomas H., John Rushing, Shalini Kansal, Sara J. Graves, Heggere Ranganath and Evans Criswell, "Eureka Phenomena Discovery and Phenomena Mining System," *Proceedings: 13th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Ocenography and Hydrology*, Long Beach, California, February, 1997.

[Hinke 1997b] Hinke, Thomas H., John Rushing, Shalini Kansal, Sara J. Graves and Heggere Ranganath. "For Scientific Data Discovery: Why Can't the Arhive be More Like the Web," *Proceedings Ninth International Conference on Scientific Database Management*, Evergreen State College, Olympia, Washington, August, 1997.

[Mining Workshop 1999] NASA Workshop on the Issues in the Application of Data Mining to Scientific Data. 1999, NASA Goddard Space Flight Center: http://www.cs.uah.edu/NASA_Mining/DMFinalReport.pdf.

[NRC 2003] Government Data Centers: Meeting Increasing Demands. 2003, National Research Council of the National Academies, Washington, DC, http://www.nap.edu.